

Page 23, lines 3-4:

~~FIG. 5 (to be illustrated later) depicts~~ FIGS. 5A and 5B depict constrained reinforcement learning algorithm in accordance with a preferred embodiment of the current invention.

IN THE CLAIMS:

Please cancel claims 1-28.

Please add the following claims in the application:

29. (New) A method for dynamically developing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the method comprising the steps of:

- a. generating a plurality of marketing strategies;
- b. determining an optimal marketing strategy based on a state of a customer and constraints corresponding to marketing channels;
- c. deploying the determined optimal marketing strategy;
- d. recording customer response to the deployed optimal marketing strategy;
- e. updating information corresponding to the state of a customer based on the recorded customer response; and
- f. repeating steps b to e for the specified time period.

30. (New) The method as recited in claim 29 wherein the step of generating a plurality of marketing strategies comprises the steps of:

selecting at least one initiative that enables an addressing of the specified objective;

determining sequences in which selected initiatives can be deployed, if more than one initiative is selected; and

combining the selected initiatives in the determined sequences to obtain the plurality of marketing strategies.

31. (New) The method as recited in claim 30 further comprising varying parameters of initiatives to generate new initiatives.

32. (New) The method as recited in claim 30 further comprising varying deployment time of initiatives.

33. (New) The method as recited in claim 29 wherein the step of determining an optimal marketing strategy further comprises the steps of:

determining all possible states of customers;

determining an optimal policy for each state based on past data;

identifying the state of a customer, the customer visiting a merchant or the customer being selected from a database of customers; and

identifying an optimal marketing strategy using the state of the customer, the identified optimal policy and constraints corresponding to marketing channels.

34. (New) The method as recited in claim 33 wherein the step of identifying all possible states of customers comprises the steps of:

identifying all relevant attributes of customers; and

partitioning the customers into partitions based on identified attributes using a similarity measure based on a historic policy, actual rewards and transition probabilities from one data point to another, the partitions forming new states of the customers.

35. (New) The method as recited in claim 33 wherein the step of determining the optimal policy for each state based on past data comprises the steps of:

identifying a deterministic policy;

initializing a value of all possible states for the policy;

computing the value of a state for the policy;
repeating said step of computing for all possible states;
constructing a new improved policy;
iteratively performing steps of computing, repeating, and constructing until the new improved policy remains unchanged for two subsequent iterations; and
selecting the policy with maximum value for the state as the optimal policy for the given state.

36. (New) The method as recited in claim 35 wherein the step of computing the value of a state for the policy comprises the steps of:

computing transition probabilities from a given state to another state for the policy;
computing value of expected immediate reward for the policy in the state;
computing discounted expected value of a resulting state for the policy; and
computing a sum of expected immediate reward and the discounted expected value.

37. (New) The method as recited in claim 35 wherein the step of constructing a new improved policy comprises the steps of:

selecting the marketing strategy which maximizes a value for the state over all marketing strategies for a given state; and
repeating said step of selecting for each state.

38. (New) The method as recited in claim 33 wherein the step of identifying an optimal marketing strategy comprises the steps of:

identifying the optimal policy for an identified customer state;
modeling customer's preferences for marketing channels, cost and effectiveness of different marketing channels, and the specified budget as effective constraints;
determining an optimal feasible policy based on the identified optimal policy and effective constraints corresponding to marketing channels; and
determining the optimal marketing strategy from the optimal feasible policy.

39. (New) The method as recited in claim 38 wherein the step of determining an optimal feasible policy based on effective constraints corresponding to marketing channels comprises mapping the optimal policy uniquely to a closest feasible optimal policy based on the effective constraints, if the effective constraints are not satisfied by the optimal policy.

40. (New) The method as recited in claim 29 wherein the step of updating information corresponding to the state of a customer based on the recorded customer response comprises the steps of:

- identifying a resulting state of the customer;
- updating values of the state of the customer; and
- updating an optimal policy.

41. (New) The method as recited in claim 40 wherein the step of updating the values of the state of the customer comprises:

- computing a sum of a new immediate reward, a discounted value corresponding to the resulting state, reduced by a value corresponding to an initial state of the customer;
- updating the values corresponding to the initial state of the customer by adding a fraction of the computed sum to a value of a previous state of the customer; and
- propagating a change in the value of the state to all other states.

42. (New) The method as recited in claim 40 wherein the step of updating the optimal policy comprises:

- computing a sum of a new immediate reward, a discounted value corresponding to the resulting state, reduced by a value corresponding to an initial state of the customer;
- and

- updating the optimal policy corresponding to an initial state of the customer by adding a fraction of the computed sum to the value of a previous state of the customer.

43. (New) A system for dynamically developing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the system comprising:
- a generator operable for generating a plurality of marketing strategies;
 - a first unit operable for determining an optimal marketing strategy based on state of a customer and constraints corresponding to marketing channels;
 - a second unit operable for deploying the determined optimal marketing strategy;
 - a recorder operable for recording customer response to the deployed optimal marketing strategy; and
 - a third unit operable for updating information corresponding to the state of a customer based on the recorded customer response.
44. (New) The system as recited in claim 43 wherein said generator comprises:
- a selector operable for selecting at least one initiative that enables an addressing of the specified objective;
 - a first sub-unit operable for determining sequences in which selected initiatives can be deployed, if more than one initiative is selected; and
 - a second sub-unit for combining the selected initiatives in the determined sequences to obtain the plurality of marketing strategies.
45. (New) The system as recited in claim 43 wherein the first unit comprises:
- a first sub-unit operable for determining all possible states of customers;
 - a second sub-unit operable for determining an optimal policy for each state based on past data;
 - a third sub-unit operable for identifying the state of a customer, the customer visiting a merchant or the customer being selected from a database of customers;
 - a fourth sub-unit operable for identifying the optimal policy for an identified customer state;

a fifth sub-unit operable for modeling customer's preferences for marketing channels, cost and effectiveness of different marketing channels, and the specified budget as effective constraints;

a sixth sub-unit operable for determining an optimal feasible policy based on effective constraints corresponding to marketing channels; and

a seventh sub-unit operable for determining the optimal marketing strategy from the optimal feasible policy.

46. (New) The system as recited in claim 45 wherein the second sub-unit comprises:

a first component operable for identifying a deterministic policy;

a second component operable for initializing a value of all possible states for the policy;

a third component operable for computing the value of a state for the policy;

a fourth component operable for constructing a new improved policy;

a fifth component operable for iteratively implementing said third component and said fourth component; and

a sixth component operable for selecting the policy with maximum value for the state as the optimal policy for the given state.

47. (New) The system as recited in claim 46 wherein the fourth component comprises a selector operable for selecting the marketing strategy that maximizes a value for the state over all marketing strategies for a given state.

48. (New) The system as recited in claim 43 wherein the third unit comprises:

a first sub-unit operable for identifying a resulting state of the customer;

a second sub-unit operable for updating a values of the state of the customer; and

a third sub-unit operable for updating an optimal policy.

49. (New) A program storage device readable by computer, tangibly embodying a program of instructions executable by the computer to perform a method for dynamically developing a marketing strategy to address at least one specified merchant objective, the

objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the method comprising:

- generating a plurality of marketing strategies;
- determining an optimal marketing strategy based on state of a customer and constraints corresponding to marketing channels;
- deploying the determined optimal marketing strategy;
- recording customer response to the deployed optimal marketing strategy; and
- updating information corresponding to the state of a customer based on the recorded customer response.

50. (New) The program storage device as recited in claim 49 wherein the step of generating a plurality of marketing strategies comprises:

- selecting at least one initiative that enables an addressing of the specified objective;
- determining sequences in which selected initiatives can be deployed, if more than one initiative is selected; and
- combining the selected initiatives in the determined sequences to obtain the plurality of marketing strategies.

51. (New) The program storage device as recited in claim 49 wherein the step of determining an optimal marketing strategy comprises:

- determining all possible states of customers;
- determining an optimal policy for each state based on past data;
- identifying the state of a customer, the customer visiting a merchant or the customer being selected from a database of customers;
- identifying the optimal policy for an identified customer state;
- modeling customer's preferences for marketing channels, cost and effectiveness of different marketing channels, and the specified budget as effective constraints;
- determining an optimal feasible policy based on effective constraints corresponding to marketing channels; and

determining the optimal marketing strategy from the optimal feasible policy.

52. (New) The program storage device as recited in claim 51 wherein the step of determining the optimal policy for each state based on past data comprises:

- identifying a deterministic policy;

- initializing a value of all possible states for the policy;

- computing the value of a state for the policy;

- constructing a new improved policy;

- iteratively executing said steps of computing and constructing; and

- selecting the policy with maximum value for the state as the optimal policy for the given state.

53. (New) The program storage device as recited in claim 52 wherein the step of constructing a new improved policy comprises selecting the marketing strategy that maximizes a value for the state over all marketing strategies for a given state.

54. (New) The program storage device as recited in claim 49 wherein the step of updating information corresponding to the state of a customer based on the recorded customer response comprises:

- identifying a resulting state of the customer;

- updating values of the state of the customer; and

- updating an optimal policy.

55. (New) A system suitable for developing an optimal marketing strategy, the system comprising:

- a database storing information regarding initiatives that can be offered to customers, marketing channels available for executing the initiatives, cost and effectiveness of the marketing channels, and states of customers;

- a unit operable for enabling a merchant to specify at least one objective for a specified time period;

a generator operable for generating a plurality of marketing strategies based on the objective specified by the merchant, the marketing strategies being a combination of initiatives; and

a component operable for determining the optimal marketing strategy and at least one marketing channel based on a state of a customer and cost and effectiveness of marketing channels.

56. (New) A method for dynamically developing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the method comprising the steps of:

- a. generating a plurality of marketing strategies;
- b. determining all possible states of customers;
- c. determining an optimal policy for each state based on past data;
- d. identifying the state of a customer, the customer visiting a merchant or the customer being selected from a database of customers;
- e. identifying the optimal policy for an identified customer state;
- f. modeling customer's preferences for marketing channels, cost and effectiveness of different marketing channels, and the specified budget as effective constraints;
- g. determining an optimal feasible policy based on the identified optimal policy and effective constraints corresponding to marketing channels;
- h. determining an optimal marketing strategy from the optimal feasible policy;
- i. deploying the determined optimal marketing strategy;
- j. recording customer response to the deployed marketing strategy;
- k. identifying a resulting state of the customer;
- l. updating values of the state of the customer;
- m. updating the optimal policy; and
- n. repeating steps c to m for the specified time period.